SAP

HA and DR Guide

Issue	01
Date	2019-07-12

HUAWEI

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Contents

1 Purpose

This document describes how to configure high availability (HA) and disaster recovery (DR) for the SAP system deployed on the public cloud. HA and DR modes can prevent single points of failure (SPOFs), greatly improving the reliability of the SAP system.

Note that this document cannot replace standard SAP installation documents. For more information about SAP deployment on the public cloud and public cloud technical components, see the operation guides of SAP and Huawei.

This document describes how to set up and con igure the system replication for SAP HANA databases, but does not guide how to con igure the high availability (HA) for databases in details, for example, using SQL Server Mirroring or Oracle Data Guard. For details, see HA installation and con iguration guides of the corresponding database.

This document does not cover all service reasons and criteria for ensuring HA on the public cloud.

We assume that you have the following knowledge or skills:

- Basic SAP concepts, operations of SAP applications and the SAP HANA database, and deployment and supported DR solutions
- HUAWEI CLOUD ECS, EVS, IMS, SFS, OBS, SDRS, CSBS, and VBS. If you are new to HUAWEI CLOUD, log in to the HUAWEI CLOUD Help Center (**https://support-intl.huaweicloud.com/en-us/index.html**) to obtain details about the cloud services.
- Knowledge about SAP system administrators
- Knowledge about Linux operating systems
- SAP applications and the HA and DR solutions for SAP HANA
- HA principles of different databases
- Knowledge about SAP HANA database administrators
- HA principles of the public cloud

2 Advantages of Cloud-Based HA

Many enterprises have deployed SAP HANA on the public cloud to run their applications. For the continuity of SAP services, the high availability (HA) and disaster recovery (DR) feature provided by the SAP on Cloud solution is important to enterprises' cloud migration. In addition to native SAP HA support, the public cloud HA and DR solution further enhances the stability and performance of SAP and many other applications. Compared with traditional HA, cloud-based HA has the following advantages:

- Agility: HA systems can be deployed on the cloud quickly and effectively. However, to realize traditional HA, you need to purchase servers and hardware, complete hardware planning, and require skilled engineers for construction.
- Flexible expansion: Cloud-based HA systems can be dynamically expanded as required to meet system performance requirements. However, the traditional HA architecture is fixed and not easy to be dynamically expanded.
- O&M: Tenants do not need to maintain the cloud-based HA. It is maintained by the cloud service providers. Traditional HA requires a local O&M team.
- Reliability: Cloud-HA can be deployed in different regions based on the tenant's requirements to achieve cross-region HA and ensure system security.
- Cost: Compared with traditional HA, cloud-based HA is cost-effective.

3 Introduction to Public Cloud Services

Public cloud refers to a cloud based on the standard cloud computing model, in which service providers make resources available to the general public over the Internet. HUAWEI CLOUD is a public cloud, and the following cloud services are required to implement SAP HANA HA.

Table 3-1 Public cloud services

Service	Description
Virtual Private Cloud (VPC)	VPC enables you to create private, isolated virtual networks. You can configure IP address segments, subnets, and security groups, assign EIPs, and allocate bandwidth in a VPC. **NOTE** All SAP HANA systems in a suit are deployed in the same VPC.
Image Management Service (IMS)	IMS provides comprehensive image management capabilities. You can select desired images from Marketplace Image to create ECSs in batches. **NOTE** Currently, the public cloud platform supports the following SAP images: • SUSE Linux Enterprise Server (SLES) 12 SP3 for SAP • SUSE Linux Enterprise Server (SLES) 12 SP4 for SAP • SUSE Linux Enterprise Server (SLES) 12 SP5 for SAP • SUSE Linux Enterprise Server (SLES) 15 SP1 for SAP
Elastic Cloud Server (ECS)	ECS provides scalable, on-demand cloud servers for secure, flexible, and efficient application environments, ensuring reliable, uninterrupted services. **NOTE** Huawei public cloud provides HANA ECSs for running SAP HANA. For more information, see the *SAP HANA User Guide*.

Service	Description
Elastic Volume Service (EVS)	EVS provides persistent block storage for services such as ECS and Bare Metal Server (BMS). With advanced data redundancy and cache acceleration capabilities, EVS offers high availability and durability with an extremely low latency. You need to perform operations such as initializing an EVS disk and creating a file system. **NOTE** For details about the EVS specifications in the SAP HANA scenario, see the *SAP HANA User Guide*.
(Optional) Scalable File Service (SFS)	SFS provides completely hosted sharable file storage for ECSs. Compatible with the Network File System protocol, SFS is expandable to petabytes, features high performance, and seamlessly handles data-intensive and bandwidth-intensive applications. **NOTE** The shared and backup volumes can be provided by SFS. For details, see the *SAP HANA User Guide*.
Volume Backup Service (VBS)	Volume Backup Service (VBS) provides online backup for EVS disks and allows data restoration from backups, offering protection against virus attacks and software and hardware faults.
Object Storage Service (OBS)	OBS is a stable, secure, efficient, and easy-to-use cloud storage service. With Representational State Transfer Application Programming Interfaces (REST APIs), OBS is able to store unstructured data of any amount and form at 99.999999999% reliability (11 nines).

4 Concepts Related to HA and Disaster Recovery

- **Availability**

Availability is used to measure the running continuity of the system. It is expressed in percentage and is inversely proportional to the downtime. For example, if the availability of a system is 99.9%, its annual downtime must be less than 0.1% or 9 hours.

- **Downtime**

Downtime is the result of service interruption, which may be caused by system upgrade or accident faults. A fault may be a device, software, or network fault, or a major disaster, such as a fire. A power failure or construction accident may cause the failure of the entire data center.

- **HA**

HA is ensured by a set of business technologies, engineering practices, and design principles. System continuity can be achieved by eliminating a single point of failure. HA minimizes the service loss by restoring the system immediately after it is interrupted.

- **DR**

Disaster recovery is performed after a data center is interrupted or onsite faults are being rectified. Data backups prepared in event of disasters are more complex and expensive.

- **Data synchronization mode**

  The SAP HANA System Replication supports synchronization (full synchronization, memory synchronization, and synchronization) and asynchronization modes. In the HA automatic switchover scenario, the synchronization mode is recommended to ensure that no data is lost during the data switchover (RPO = 0).

  – **Full synchronization**: After the secondary system receives the synchronization data and stores it in disks, the secondary system sends confirmation information to the primary system and then the primary system submits logs of this operation. If the secondary system fails to

receive the synchronization data due to a fault, the primary system waits until the secondary system recovers.

- **Synchronization**: After the secondary system receives the synchronization data and stores it in disks, the secondary system sends confirmation information to the primary system and then the primary system submits logs of this operation. If the secondary system fails to receive synchronization data due to a fault and the primary system fails to synchronize data to the secondary system, the primary system continues its service.
- **Memory synchronization**: After the secondary system receives the logs (memory) and sends a confirmation message to the primary system, the primary system submits logs of this operation. If the secondary system fails to receive synchronization data due to a fault and the primary system fails to synchronize data to the secondary system, the primary system continues its service.
- **Asynchronization**: The primary system sends the synchronization data to the secondary system and submits logs without waiting for the response from the secondary system. The fault of the secondary system does not affect the primary system running services.

- **Data preload mode**
  - The preload option is **ON**.

    Data is stored in the memory.

    The switchover is fast (RTO).
  - The preload option is **OFF**.

    The memory usage is low and the memory is used by other systems, such as non-production systems.

    The switchover duration is long.

    Before the switchover, activate the preload mode.

5 SAP HANA HA and DR

5.1 Native SAP HANA HA Support

SAP provides a set of native HA support to recover the SAP HANA system from various faults.

Service Auto-Restart

In the event of a software failure or an intentional intervention by an administrator that disables one of SAP HANA services, the service will be restarted by the SAP HANA service auto-restart watchdog function, which automatically detects the failure and restarts the stopped service process. Upon restart, the service loads data into memory and resumes its function. SAP HANA service auto-restart can be used right out of the box for fault recovery, which works the same way on the HUAWEI CLOUD as on other platforms.

Host Auto-Failover

This solution requires that SAP HANA be installed on multiple nodes (at least two), including one master node, several slave nodes, and one or more standby nodes. When a single point of failure occurs on the primary node or slave node, the standby node can identify and automatically switch from the standby state to the active state to replace the failed node.

HANA System Replication (HSR)

System replication is available in every SAP HANA installation offering HA and DR support. A secondary system is set up to copy data and transaction logs from the primary system. Once a fault occurs in the primary system, the secondary system is enabled to take over operations from the primary system. In addition, HSR allows you to have a variety of settings, depending on your requirements for recovery time objective (RTO) and database recovery point objective (RPO). For details, see **How to Perform System Replication for SAP HANA** released by SAP.

Backup and Recovery

Although SAP HANA is an in-memory database, all data modifications are saved in a persist storage system. Therefore, even if a sudden power outage occurs, SAP

HANA can restore data and ensure data integrity. In addition, to restore data after a disaster occurs, data and database logs in the storage system must be backed up to a remote end, for example, OBS bucket. For more details about SAP HANA database backup and recovery, see **Backup and Recovery - SAP HANA** released by SAP.

5.2 Public Cloud HA and DR Scenarios

5.2.1 Automatic Recovery of SAP HANA ECSs

Scenarios

The public cloud provides automatic recovery by default to restart ECSs through cold migration, ensuring high availability and dynamic ECS migration.

- You can enable automatic recovery during or after ECS creation.
- An ECS can be automatically recovered only if the host on which it is deployed becomes faulty. This function does not take effect if the fault is caused by the ECS itself.
- An ECS can be automatically recovered only once within 12 hours if the host on which it is deployed becomes faulty.
- ECS automatic recovery may fail in the following scenarios:
  - Many hosts in the system are faulty.
  - The host to which the ECS services are migrated does not have sufficient temporary capacity.

Advantages

The public cloud provides the automatic recovery capability for ECSs with easy configuration and no extra cost.

NOTICE

It is recommended to use the automatic recovery capability of ECSs in the SAP non-production environment.

Deployment Plan

Figure 5-1 shows the deployment plan.

Figure 5-1 Automatic recovery of SAP HANA ECSs

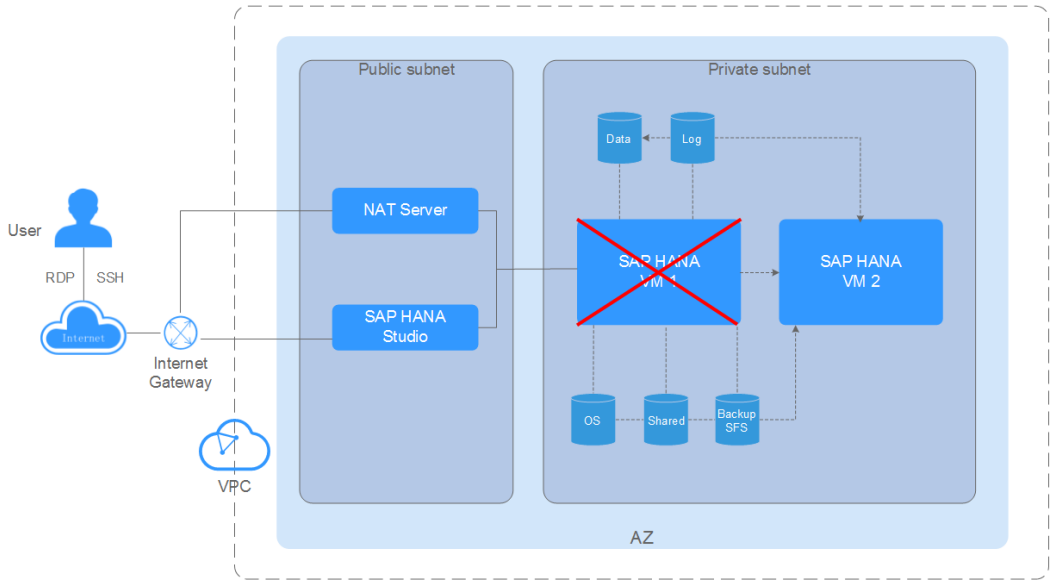

5.2.2 SAP HANA Cluster HA (Host Auto-Failover)

Scenarios

The public cloud supports the SAP HANA cluster mode, which has the HA capability ensured by the configured standby nodes and the automatic switchover function (Host Auto-Failover).

- Standby nodes are configured for the SAP HANA cluster to implement automatic switchover. When any node in the cluster fails, the SAP HANA automatically triggers a switchover. The switchover does not cause data loss (RPO = 0).
- A single standby node can only recover the fault of a single node at one time. It is recommended to configure multiple standby nodes for an SPA HANA cluster with over eight nodes.
- The public cloud supports anti-affinity deployment of the SAP HANA cluster. ECSs in the same SAP HANA cluster are deployed on different physical servers to ensure HA.
- Currently, the public cloud supports a maximum of eight nodes (7 worker nodes and 1 standby node) in an SAP HANA cluster.

Advantages

The SAP HANA cluster HA is easy to be deployed and cost-effective based on its HA mechanism and Huawei public cloud capabilities.

NOTICE

Before enabling the HA automatic switchover function of an SAP HANA cluster, disable the automatic recovery function of the ECSs in the cluster.

Deployment Plan

Figure 5-2 shows the deployment plan.

Figure 5-2 SAP HANA cluster HA (Host Auto-Failover)

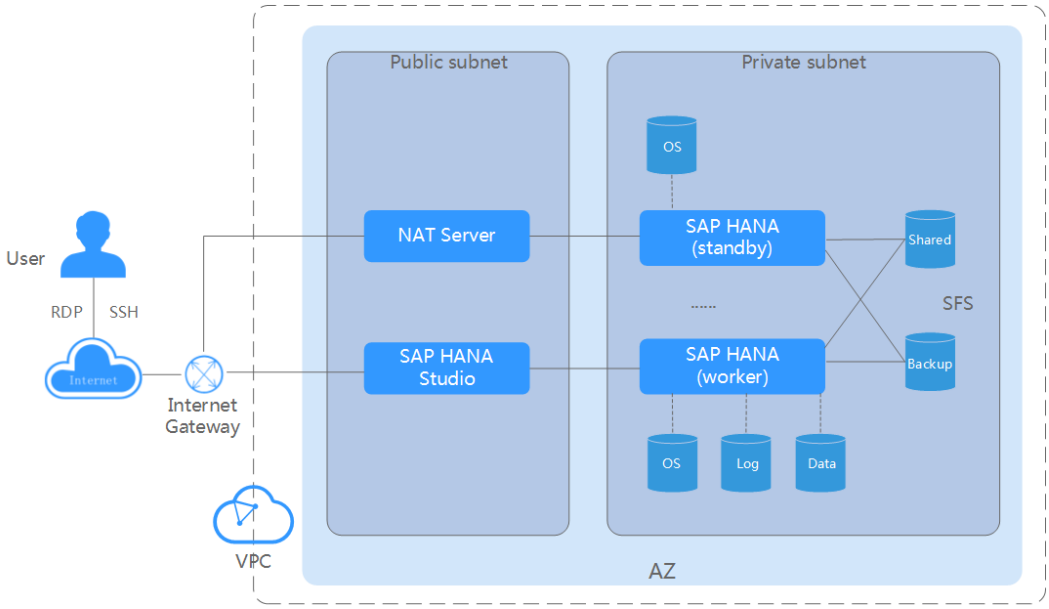

5.2.3 SAP HANA Single-Node HA (System Replication + Preload on + SUSE HAE)

SAP HANA HA within an AZ

Scenarios

The public cloud supports SAP HANA HA. In this scenario, deploy the primary and secondary SAP HANA systems in the same AZ and configure the System Replication function of SAP HANA and the SUSE High-Availability Extension (HAE) function to enable automatic switchover and realize HA. This applies to the scenario where a single SAP HANA node is deployed.

- The System Replication function of SAP HANA provides data replication.
- Enabling the Preload function of System Replication can preload data to the memory of the secondary system to reduce the switchover time and lower Recovery Time Objective (RTO).
- Configuring SUSE HAE allows automatic system switchover. The public cloud provides scripts for automatic HAE configuration. For details, see the *SAP HANA User Guide (Single Node Deployment)*.
- The HA mechanism automatically monitors the health status of the primary SAP HANA system (SAP HANA process, OS, and data replication). When the primary system fails, failover and virtual IP address floating are automatically triggered to achieve HA of the HANA system.

Advantages

Based on the public cloud capabilities, SAP HANA System Replication, and SUSE HAE, SAP HANA provides reliable assurance for your core services.

NOTICE

In the production environment, you are advised to use the SAP HANA HA automatic switchover solution to reduce the RTO. Before using this solution, you need to disable the automatic recovery function of ECSs.

Deployment Plan

Figure 5-3 shows the deployment plan.

Figure 5-3 SAP HANA HA within an AZ

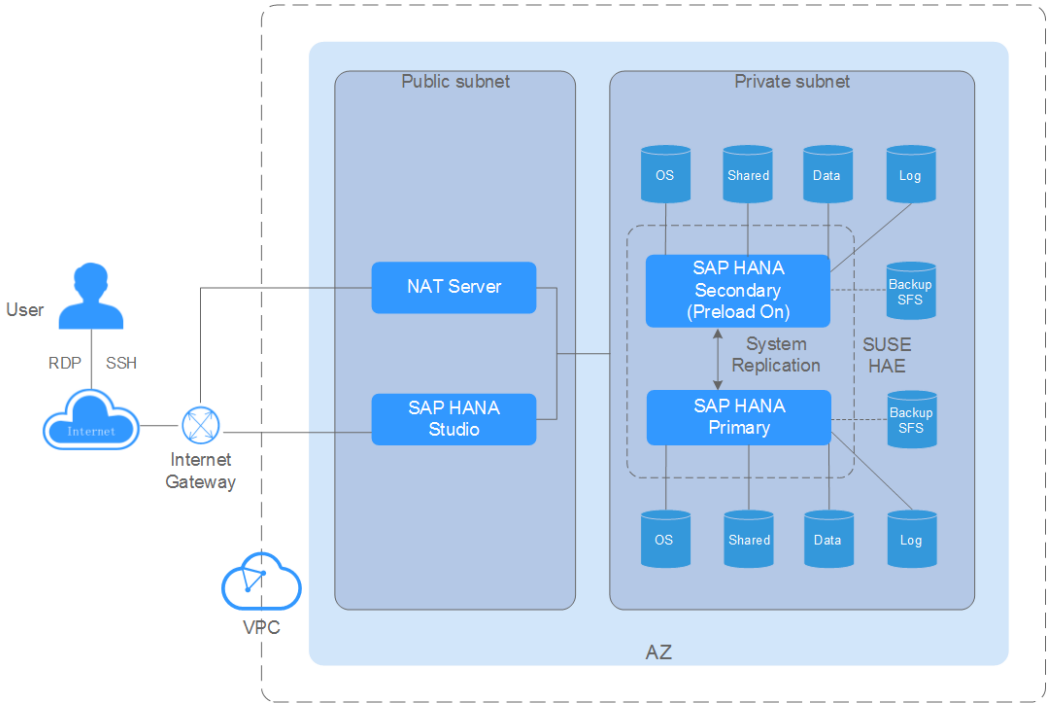

Cross-AZ SAP HANA HA

Scenarios

Deploy secondary SAP HANA systems in different AZs.

- The System Replication function of SAP HANA provides data replication.
- Enabling the Preload function of System Replication can preload data to the memory of the secondary system to reduce the switchover time and lower Recovery Time Objective (RTO).
- Configuring SUSE HAE allows automatic system switchover. The public cloud provides scripts for automatic HAE configuration. For details, see the *SAP HANA User Guide (Single Node Deployment)*.
- The HA mechanism automatically monitors the health status of the AZ where the primary system (SAP HANA process, OS, and data replication) is located.

When the AZ where the primary system is located fails, failover and virtual IP address floating are automatically triggered to achieve HA of the HANA system.

Advantages

Based on the public cloud capabilities, SAP HANA System Replication, and SUSE HAE, cross-AZ SAP HANA deployment provides reliable assurance for your core services.

Deployment Plan

Figure 5-4 shows the deployment plan.

Figure 5-4 Cross-AZ SAP HANA HA

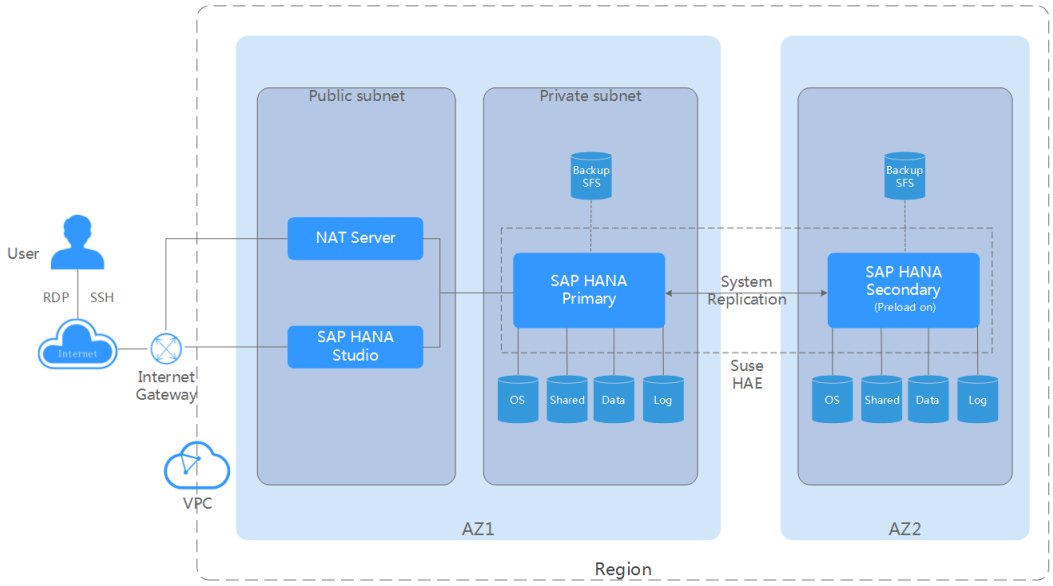

Scenario 1: SAP HANA HA (System Replication + Preload off + DEV/QAS)

Scenarios

After the SAP HANA HA is deployed, the Preload function of System Replication is disabled. In this case, the secondary system can be used to deploy non-production systems.

- After the preload function of System Replication is disabled, only a few resources (10%) need to be reserved for the secondary system for receiving synchronization data from the primary system. In addition, non-production systems can be deployed on the secondary system.
- According to the SAP best practice, non-production systems on the secondary system have to use additional storage resources to avoid the impact on the production system after the switchover of the secondary system.
- Before the switchover, disable non-production systems on the secondary system to ensure sufficient resources on the secondary system and then manually trigger the switchover.

Advantages

The secondary SAP HANA system can be fully utilized to reduce the overall cost based on public cloud capabilities and SAP HANA HA mechanism.

NOTICE

In the production environment, you are advised to use the SAP HANA single-node HA automatic switchover solution to reduce the RTO. Before using this solution, you need to disable the HA configuration function of ECSs.

Deployment Plan

Figure 5-5 shows the deployment plan.

Figure 5-5 SAP HANA HA within an AZ (System Replication + Preload off + DEV/QAS)

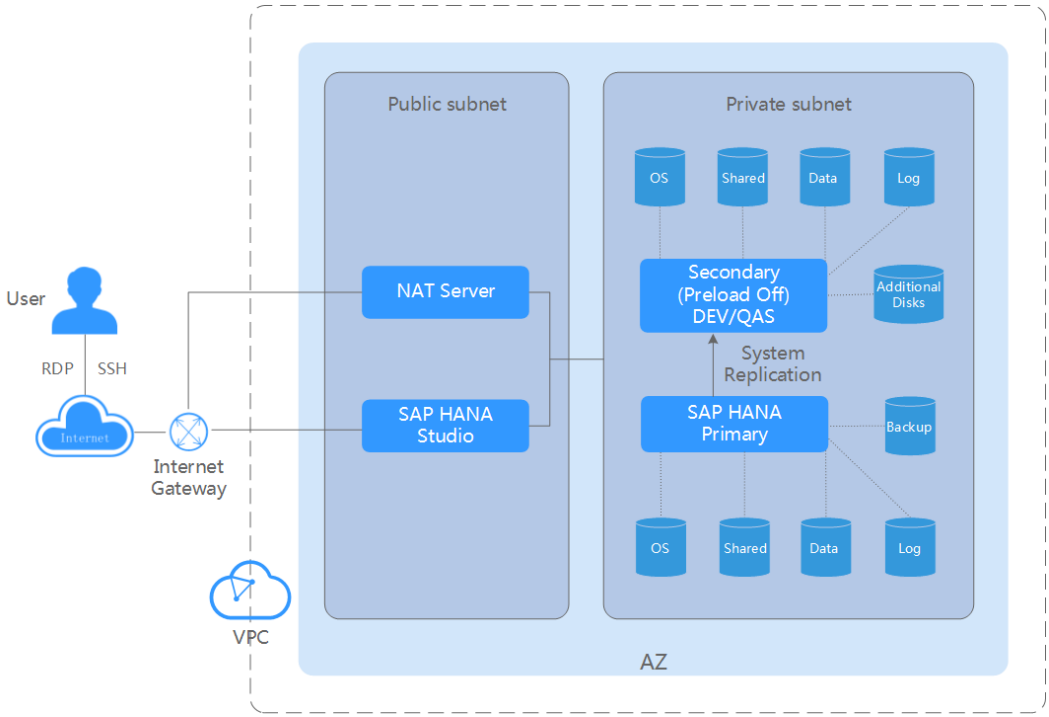

5.2.4 SAP HANA DR (System Replication)

Scenarios

The public cloud supports SAP HANA DR. Cross-AZ SAP HANA DR can be achieved by deploying primary and secondary SAP HANA systems in different AZs and using the System Replication function of SAP HANA to implement data replication.

- Based on the network status between AZs, you can select the synchronization or asynchronization mode to meet different RPO requirements.
- When the preload function is disabled, the secondary system can be used to deploy non-production systems. For details, see **Scenario 1**.

Advantages

A VPC can be deployed across AZs and allows the primary and secondary SAP HANA systems to be deployed in the same subnet. No extra network configuration is required.

Deployment Plan

Figure 5-6 shows the deployment plan.

Figure 5-6 SAP HANA DR (System Replication)

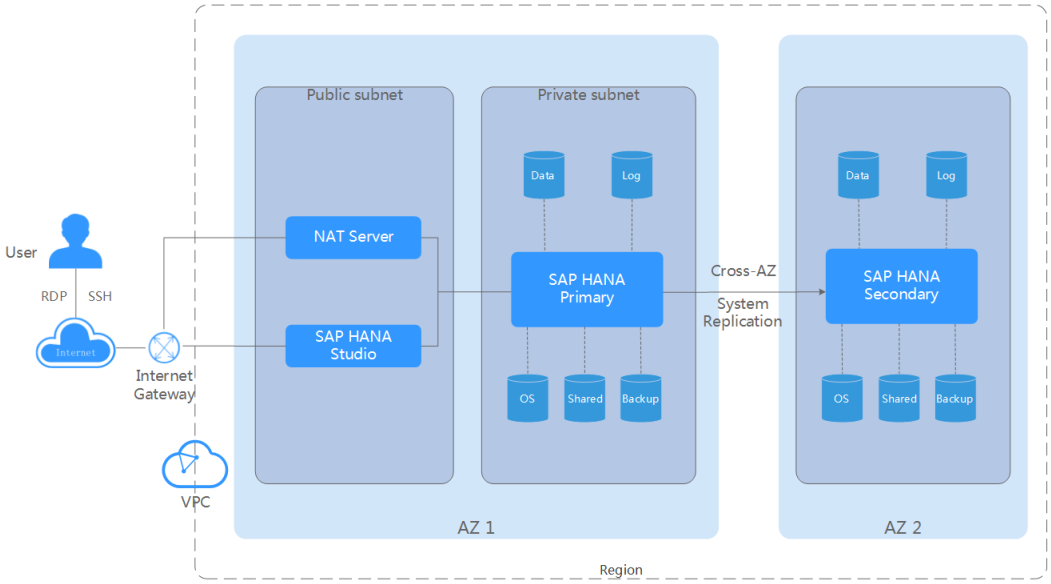

Scenario 1: SAP HANA DR (System Replication + Preload off + DEV/QAS)

The public cloud supports SAP HANA DR. Cross-AZ SAP HANA DR can be achieved by deploying primary and secondary SAP HANA systems in different AZs and using System Replication to implement data replication. The secondary system can be used to deploy non-production systems when the preload function is disabled.

- After the preload function of System Replication is disabled, only a few resources (10%) need to be reserved for the secondary system for receiving synchronization data from the primary system. In addition, non-production systems can be deployed on the secondary system.
- According to the SAP best practice, non-production systems on the secondary system have to use additional storage resources to avoid the impact on the production system after the switchover of the secondary system.
- Before the switchover, disable non-production systems on the secondary system to ensure sufficient resources on the secondary system and then manually trigger the switchover.

Advantages

A VPC can be deployed across AZs and allows the primary and secondary SAP HANA systems to be deployed in the same subnet. No extra network configuration is required. The secondary system can be fully utilized to reduce the overall cost.

Deployment Plan

Figure 5-7 shows the deployment plan.

Figure 5-7 SAP HANA DR (System Replication + Preload off + DEV/QAS)

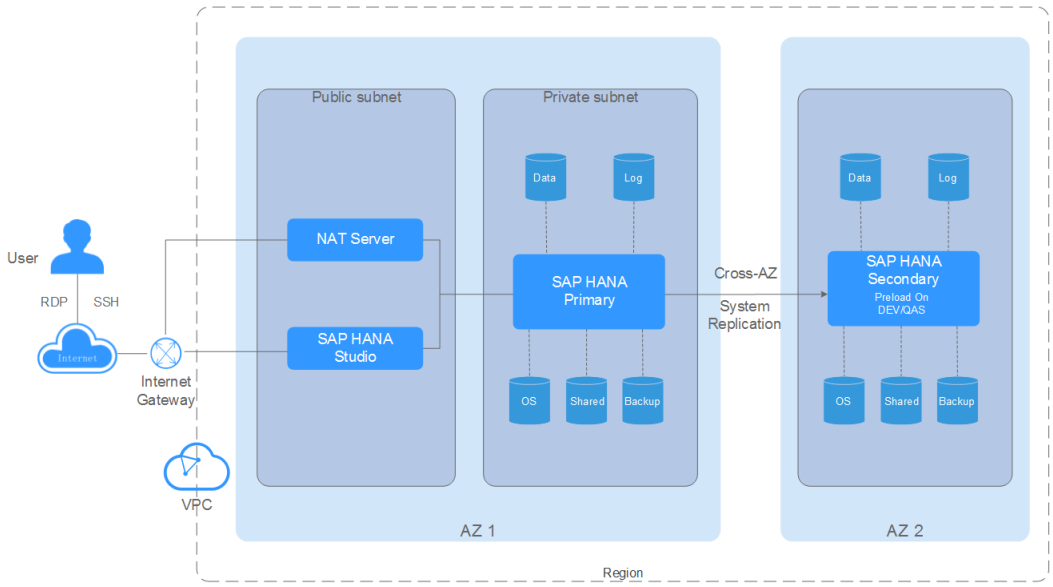

Scenario 2: Cloud DR System Deployment

The public cloud supports deploying the SAP HANA DR system on the cloud to synchronize data with the system in the customer's local data center.

- Use the VPN service to connect customer data centers to public cloud resources. Deploy the SAP HANA DR system on the cloud.
- Synchronize data between on-cloud and on-premises SAP HANA systems using System Replication in the asynchronization mode.
- Manually switch to the on-cloud SAP HANA system when a fault occurs in the customer data center system.

Advantages

Public cloud resources allow the complete DR between on-cloud and on-premises systems.

Deployment Plan

Figure 5-8 shows the deployment plan.

Figure 5-8 DR system deployment on cloud

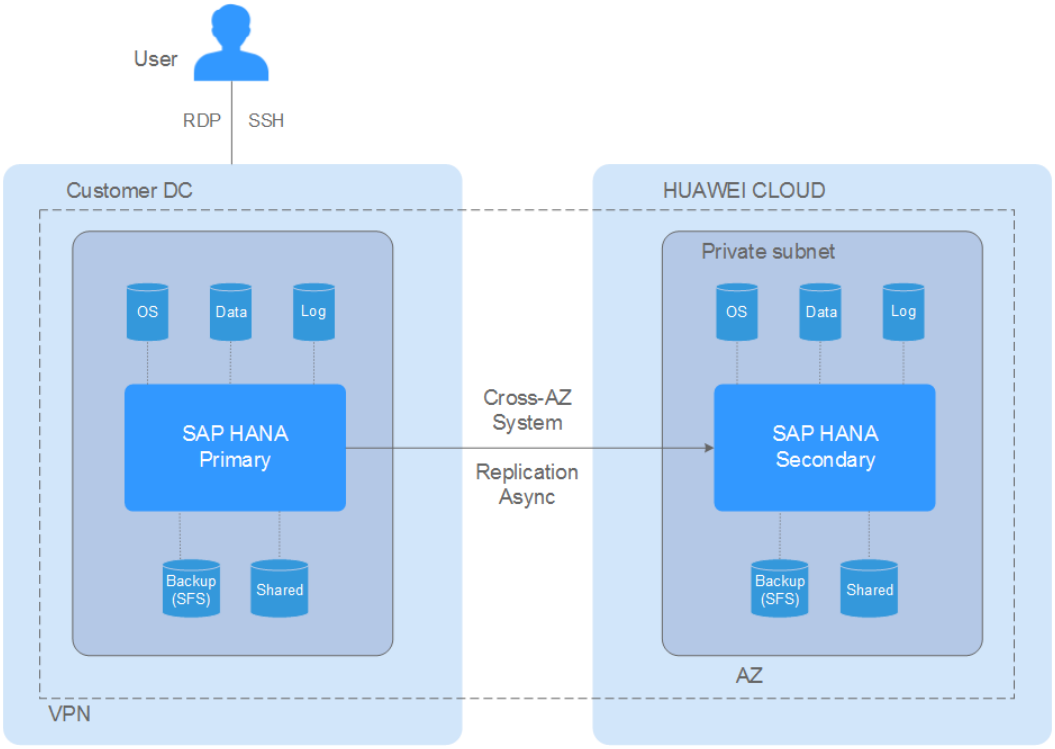

5.2.5 SAP HANA HA and DR

Scenario 1: Single-AZ SAP HANA HA (System Replication in the Synchronization Mode + Preload on + SUSE HAE) and Cross-AZ SAP HANA DR (System Replication in the Asynchronization Mode + Preload on)

Scenarios

Deploy the primary and secondary SAP HANA systems in the same AZ and configure the System Replication function of SAP HANA and the SUSE HAE function to enable automatic switchover and realize HA. In the other AZ, deploy a secondary SAP HANA system and configure Multitier System Replication function to implement cross-AZ DR. This applies to the scenario where a single SAP HANA node is deployed.

Deployment Plan

Figure 5-9 shows the deployment plan of SAP HANA HA within an AZ and SAP HANA DR across AZs.

Figure 5-9 Single-AZ SAP HANA HA and cross-AZ SAP HANA DR

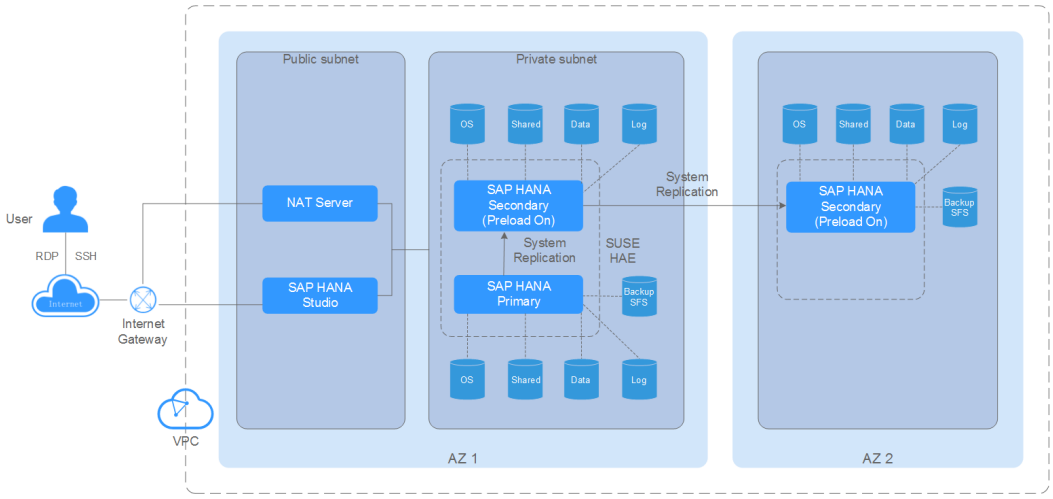

Scenario 2: Single-AZ SAP HANA HA (System Replication in the Synchronization Mode + Preload on + SUSE HAE) and Cross-Region SAP HANA DR (System Replication in the Asynchronization Mode + Preload on)

Scenarios

Deploy the primary and secondary SAP HANA systems in the same AZ and configure the System Replication function of SAP HANA and the SUSE HAE function to enable automatic switchover and realize HA. In the other region, deploy a secondary SAP HANA system and configure Multitier System Replication function to implement cross-region DR. This applies to the scenario where a single SAP HANA node is deployed.

Deployment Plan

Figure 5-10 shows the deployment plan of SAP HANA HA within an AZ and SAP HANA DR across regions.

Figure 5-10 Single-AZ SAP HANA HA and cross-region SAP HANA DR

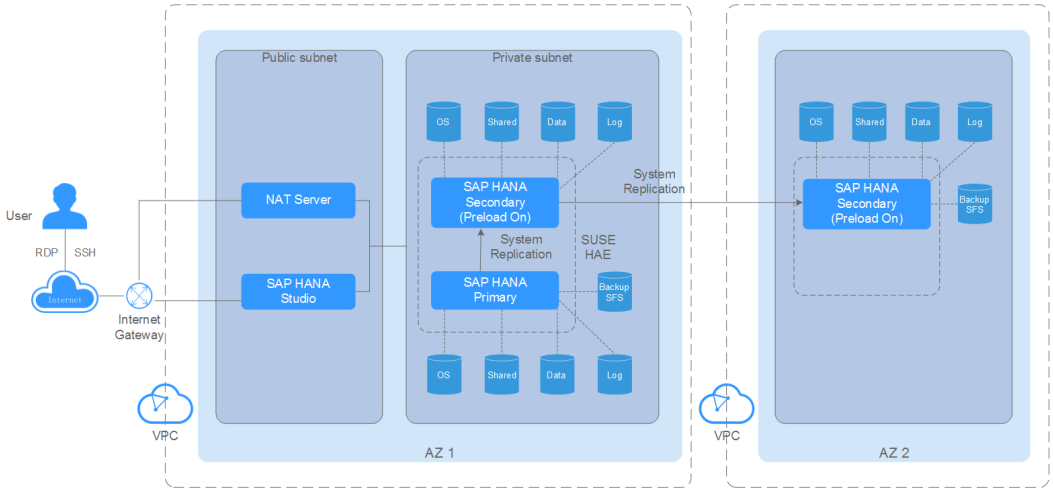

Scenario 3: Cross-AZ SAP HANA HA (System Replication in the Synchronization Mode + Preload on + SUSE HAE) and Cross-Region SAP HANA DR (System Replication in the Asynchronization Mode + Preload on)

Scenarios

Deploy the primary and secondary SAP HANA systems in different AZs of the same region and configure the System Replication function of SAP HANA and the SUSE HAE function to enable automatic switchover and realize HA across AZs. In the other region, deploy a secondary SAP HANA system and configure Multitier System Replication function to implement cross-region DR. This applies to the scenario where a single SAP HANA node is deployed.

Deployment Plan

Figure 5-11 shows the deployment plan of SAP HANA HA across AZs and SAP HANA DR across regions.

Figure 5-11 Cross-AZ SAP HANA HA and cross-region SAP HANA DR

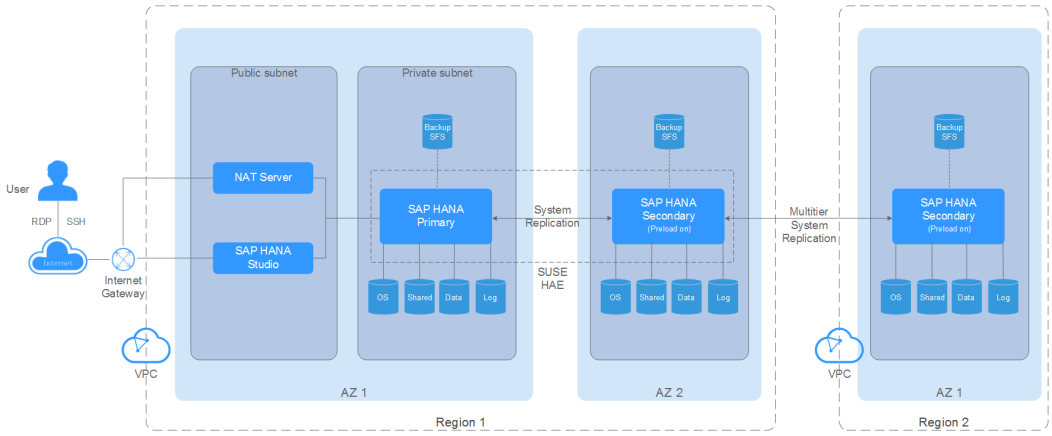

5.2.6 CSBS Application Consistency Backup and Restoration

Scenario

The public cloud backs up SAP HANA databases using the CSBS application consistency backup. In the same AZ, an SAP HANA database deployed in the single-node mode is used to store business data. As the data volume increases, the backup mode cannot meet the RTO and RPO requirements. The application consistency backup is used to reduce the RTO and RPO, maximizing user data security and correctness and ensuring business security.

- The periodic automatic backup is supported.
- The SAP HANA snapshot technology is used to back up data of multiple EVS disks contained in ECSs to the OBS.

> **NOTE**
>
> Storage snapshots cannot ensure file system consistency. The CSBS Agent integrates the SAP HANA snapshot function to ensure file system consistency.

When the CSBS application consistency backup is used, the steps in Figure 5-12 are automatically performed.

Figure 5-12 CSBS application consistency backup and restoration flowchart

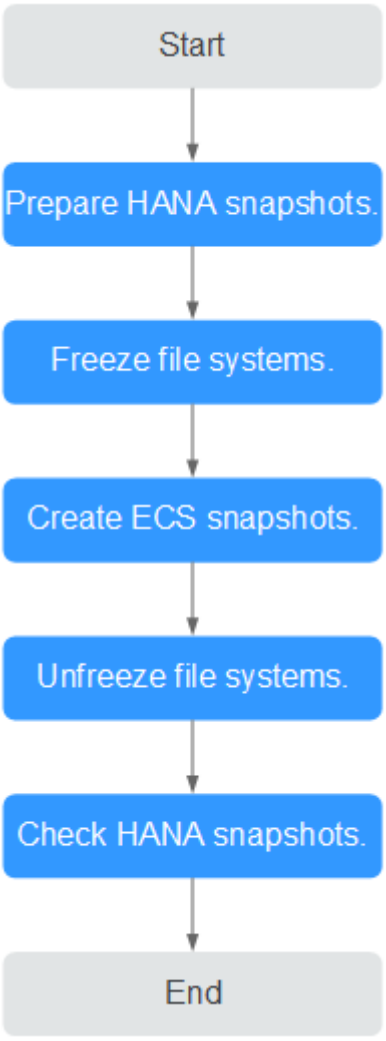

Advantages

- ECS backup and restoration: Restores the original ECS and uses backup data to create a new ECS.
- Flexible and easy to use: Users can manually trigger backup through the console or use scripts edited with the APIs provided by HUAWEI CLOUD to periodically trigger backup.
- Cross-AZ ECS restoration: Backup data is uploaded to OBS for disaster recovery at the AZ level.

Deployment Scheme

Figure 5-13 shows the deployment scheme.

Figure 5-13 Application consistency backup deployment scheme

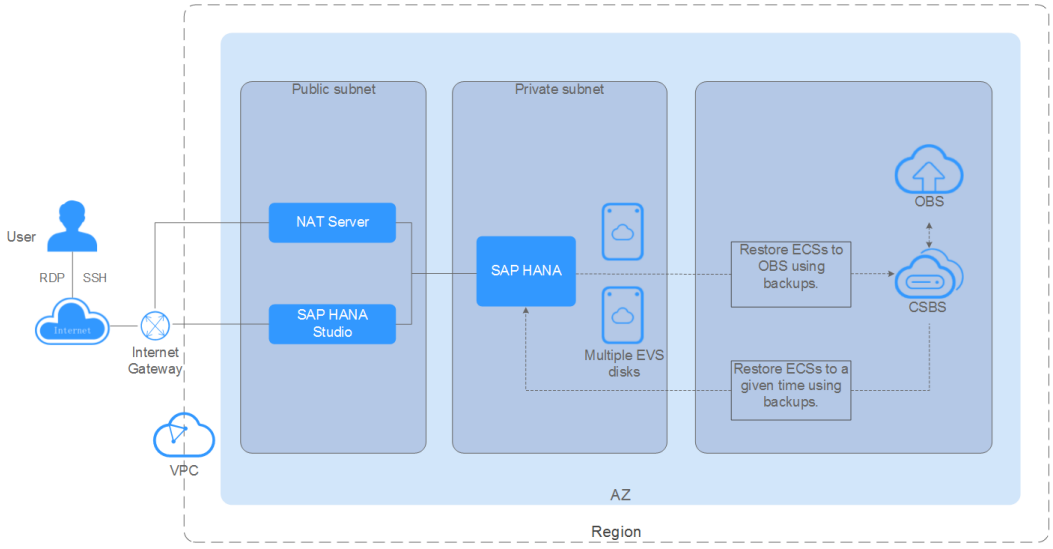

Constraints

- An ECS can be associated with only one backup policy.
- Only the application consistency backup for SAP HANA deployed in the single-node mode is supported.
- CSBS supports backup and restoration of part of the disks on an ECS as a whole but does not support file- or directory-level restoration.
- Before enabling the application consistency backup, install the required Agent software. Otherwise, the application consistency backup will fail. For details, see **Management Operations** > **Backup** > **Enabling Application-Consistency Backup** > **Installing the Agent** in the *Cloud Server Backup User Guide*.
- Application consistency backup and restoration use the ECS restoration mechanism rather than the HANA restoration mechanism, and cannot be restored using HANA snapshots.
- CSBS can only use the backup at the snapshot backup time to restore ECSs or HANA ECSs. If you want to restore an SAP HANA database at any given time, use the backup and restoration function of the SAP HANA database.

Scenario 1: Data disk restoration using the application consistency backup

Scenario

Deploy SAP HANA systems in an AZ, create an application consistency backup, create an ECS image using the backup, provision a new ECS using the image, and stop the ECS. Use the backup to create a data disk image, attach the data disk image to the newly provisioned ECS, start the ECS, and start the HANA database to complete backup and restoration. This applies to the scenario where SAP HANA is deployed in the single-node mode.

Deployment Scheme

Figure 5-14 shows the deployment scheme about how to restore a data disk using the application consistency backup when SAP HANA systems are deployed in the same AZ.

Figure 5-14 Scenario 1: Data disk restoration using the application consistency backup

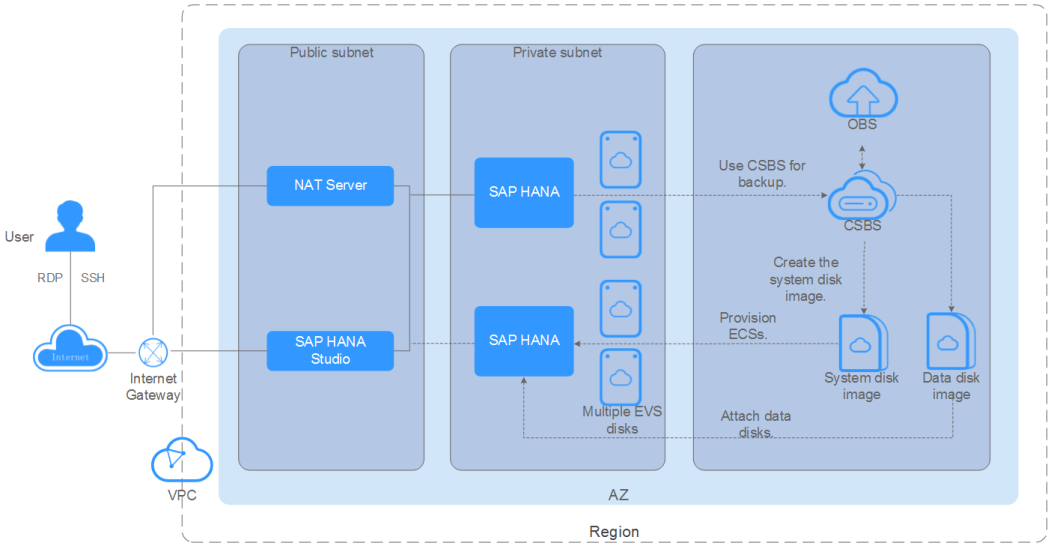

Scenario 2: Cloud server restoration using the application consistency backup

Scenarios

Deploy SAP HANA systems in the same AZ, create an application consistency backup, create an ECS image using the backup, provision an ECS using the image, change the ECS IP address, and restart HANA ECS to complete backup and restoration. This applies to the scenario where SAP HANA is deployed in the single-node mode.

Deployment Plan

Figure 5-15 shows the deployment scheme about how to create an ECS image restoration using the application consistency when SAP HANA systems are deployed in the same AZ.

Figure 5-15 Scenario 2: Cloud server image restoration using the application consistency backup

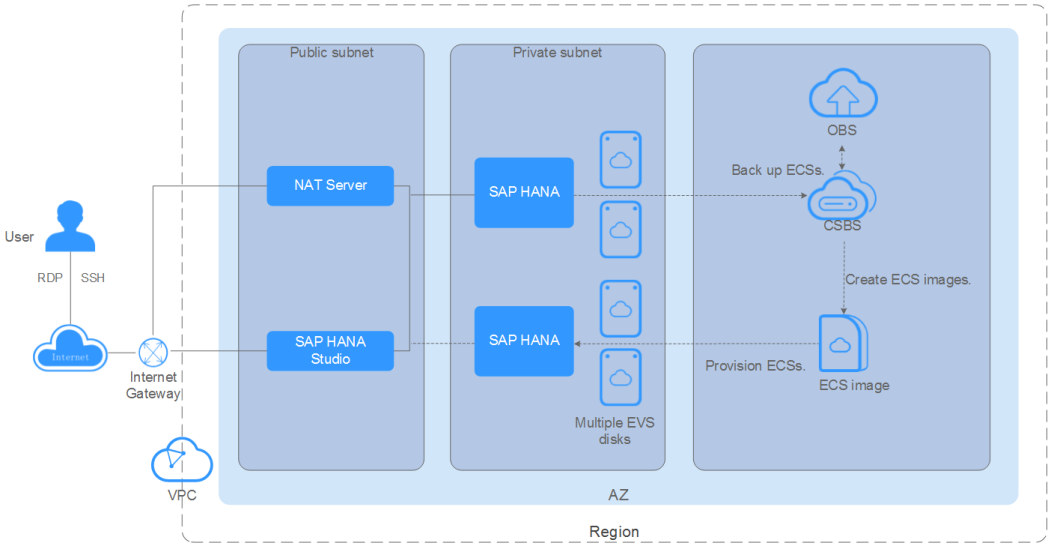

Scenario 3: Cloud server image restoration using the application consistency backup in the cross-region scenario

Scenario

Deploy SAP HANA systems in the same AZ, create an application consistency backup, copy the backup across regions, create an ECS image using the backup, provision an ECS using the image, change the ECS IP address, and restart HANA ECS to complete backup and restoration. This applies to the scenario where SAP HANA is deployed in the single-node mode.

Deployment Scheme

Figure 5-16 shows the deployment scheme about how to create an ECS image restoration using the application consistency backup across regions when SAP HANA systems are deployed in the same AZ.

Figure 5-16 Cloud server image restoration using the application consistency backup in the cross-region scenario

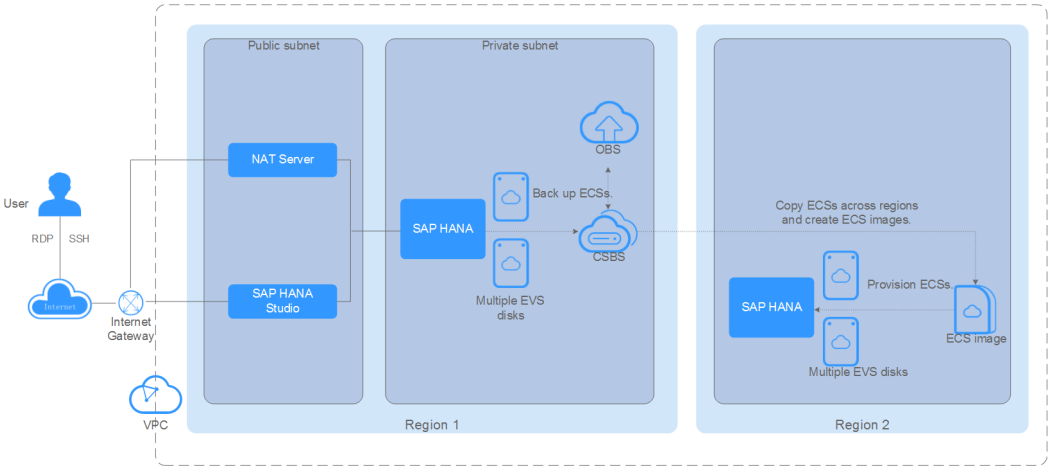

5.2.7 SAP HANA HA and DR Scenario Comparison

Scenario Comparison

Table 5-1 Scenario summary and comparison

Scenario	Deployment	Application Scope	Applicable SAP Scenario
Automatic Recovery of SAP HANA ECSs	Local deployment and HA	Single node/ Cluster	Non-production scenario (development, test, training, and sandbox)
SAP HANA HA (System Replication + Preload on+ SUSE HAE)	Local deployment and HA	Single node/ Cluster	Production scenario
SAP HANA HA (System Replication + Preload off)	Local deployment and HA	Single node/ Cluster	Production and non-production scenarios
SAP HANA DR (System Replication)	Cross-AZ deployment and DR	Single node/ Cluster	Production and non-production scenarios
On-Cloud and On-Premises SAP HANA DR (System Replication)	Cross-region deployment and DR	Single node/ Cluster	Production scenario
Application Consistency Backup and Restoration	Cross-AZ deployment and DR	Single node	Non-production scenario (development, test, training, and sandbox)

Comparison Between Cost and Reliability

Table 5-2 Comparison between cost and reliability

Scenario	Additional Cost	RPO	RTO	Remarks
Automatic Recovery of SAP HANA ECSs	None	N/A	High	RPO is uncertain. RTO is high. Restart the ECS using cold migration.

Scenario	Additional Cost	RPO	RTO	Remarks
SAP HANA HA (System Replication + Preload on+ SUSE HAE)	Medium N + N	0	Low	RPO = 0. The synchronization mode is used for data replication. RTO is in the minute level. The switchover between primary and secondary systems is automatic.
SAP HANA HA (System Replication + Preload off)	Medium N + N	≈ 0	Medium	RPO varies depending on the data replication mode and network status. RTO is high. (Decide the switchover policy. Manually perform the switchover and preload the data to the memory.)
SAP HANA DR (System Replication)	Medium to high N + 2N	≈ 0	Medium	RPO varies depending on the data replication mode and network status. RTO is high. (Decide the switchover policy. Manually perform the switchover and preload the data to the memory.)
On-Cloud and On-Premises SAP HANA DR (System Replication)	Medium N + N	> 0	Medium	RPO > zero. Generally, the asynchronization mode is used. RTO is high. (Decide the switchover policy. Manually perform the switchover and preload the data to the memory.)
CSBS Application Consistency Backup and Restoration	Low	> 0	High	RPO is uncertain. RTO is high. The application consistency backup uses the ECS restoration mechanism.

6 SAP Application HA and DR

6.1 SAP Application SDRS

6.1.1 Introduction to SDRS

HUAWEI CLOUD Storage Disaster Recovery Service (SDRS) provides cross-AZ DR protection for servers. It ensures that recovery point objective (RPO) is equal to 0, greatly reduces DR TCO for enterprises, and simplifies the DR process. If a fault occurs at the production site, you can quickly restore services at the DR site. This significantly shortens service interruptions and reduces loss.

For more information about SDRS, see **SDRS Introduction**.

NOTE

SDRS supports only cross-AZ EVS DR, and does not provide DR for file systems of the Scalable File Service (SFS) or SFS Turbo.

6.1.2 DR Solution for Standard Deployment

Table 6-1 shows the file systems used in standard SAP application deployment.

Table 6-1 File system planning

File System	Type	Cloud Service
/	ext4	EVS
/usr/sap	xfs	EVS
/sapmnt	xfs	EVS

In standard deployment mode, all file systems are mounted on EVS disks. Therefore, SDRS is applicable to the SAP application systems. **Figure 6-1** shows the DR solution for the standard deployment mode.

Figure 6-1 DR scheme

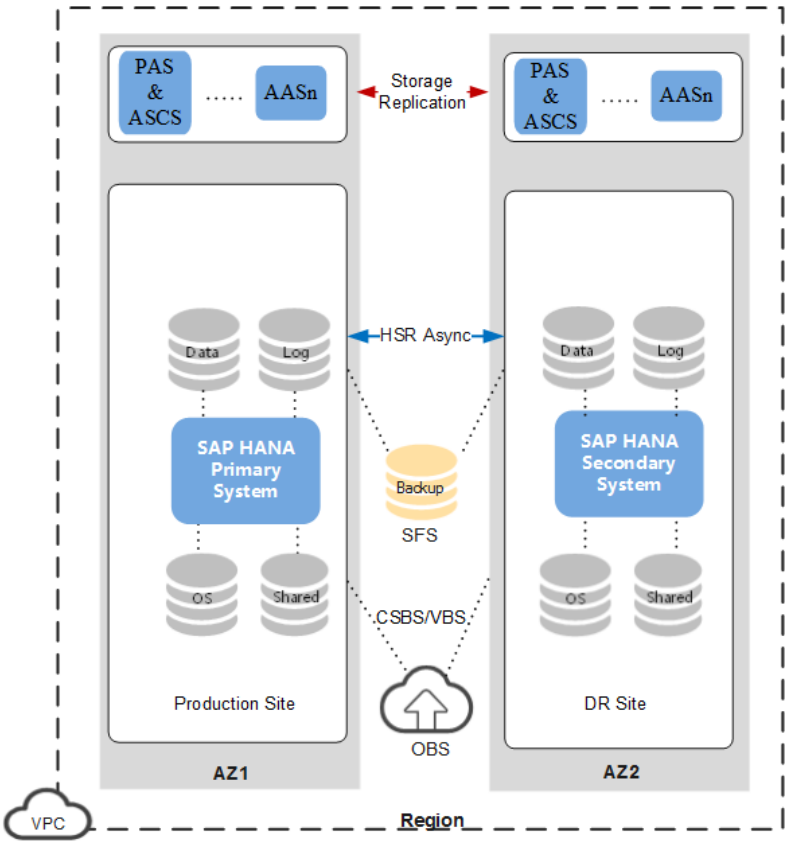

Solution Description

SAP applications are deployed in the standard mode, and SAP HANA is deployed in a single-node mode without HA required. AZ 1 is the production site, and AZ 2 is the DR site.

SDRS is configured for SAP applications, and asynchronous system replication is configured for SAP HANA. SAP applications are deployed on one PAS & ASCS node and as many AAS nodes as possible. SDRS is configured for each node. The system volume, data volume, log volume, and shared volume of SAP HANA are provided by EVS in AZ 1, and the backup volume is provided by SFS, and is mounted across AZs to both the production and DR sites. SAP HANA data is periodically backed up to the backup volume, or to an OBS bucket using CSBS or VBS.

If the production site is faulty, SAP applications are switched over to the DR site through SDRS. Servers and disks at the DR site become available immediately. SAP HANA restores to a specified data recovery point through backup files in the backup volume, or using CSBS or VBS.

6.1.3 DR Solution for Distributed Deployment

This section describes the DR solutions for SAP applications deployed in the distributed mode without HA required and the same mode with HA required.

- Table 6-2 shows the file systems used by the SAP system deployed in the distributed mode without HA required.

Table 6-2 File system planning for the distributed mode without HA

File System	Type	Cloud Service
/	ext4	EVS
/usr/sap	xfs	EVS
/sapmnt/<SID>	nfs	SFS Turbo
/usr/sap/trans	nfs	SFS Turbo

- **Table 6-3** shows the file systems used by the SAP system deployed in the distributed mode with HA required.

Table 6-3 File system planning for the distributed mode with HA

File System	Type	Cloud Service
/	ext4	EVS
/usr/sap	xfs	EVS
/usr/sap/<SID>/ASCS<##>	nfs	SFS Turbo
/usr/sap/<SID>/ERS<##>	nfs	SFS Turbo
/sapmnt	nfs	SFS Turbo
/usr/sap/<SID>/SYS	nfs	SFS Turbo

In the distributed deployment modes (with and without HA), the file systems use EVS disks and SFS. SDRS is applicable to file systems that use EVS disks. For file systems use SFS, you can back up and restore the files for DR. **Figure 6-2** shows the DR solution for SAP applications deployed in the distributed mode with HA required.

Figure 6-2 DR solution

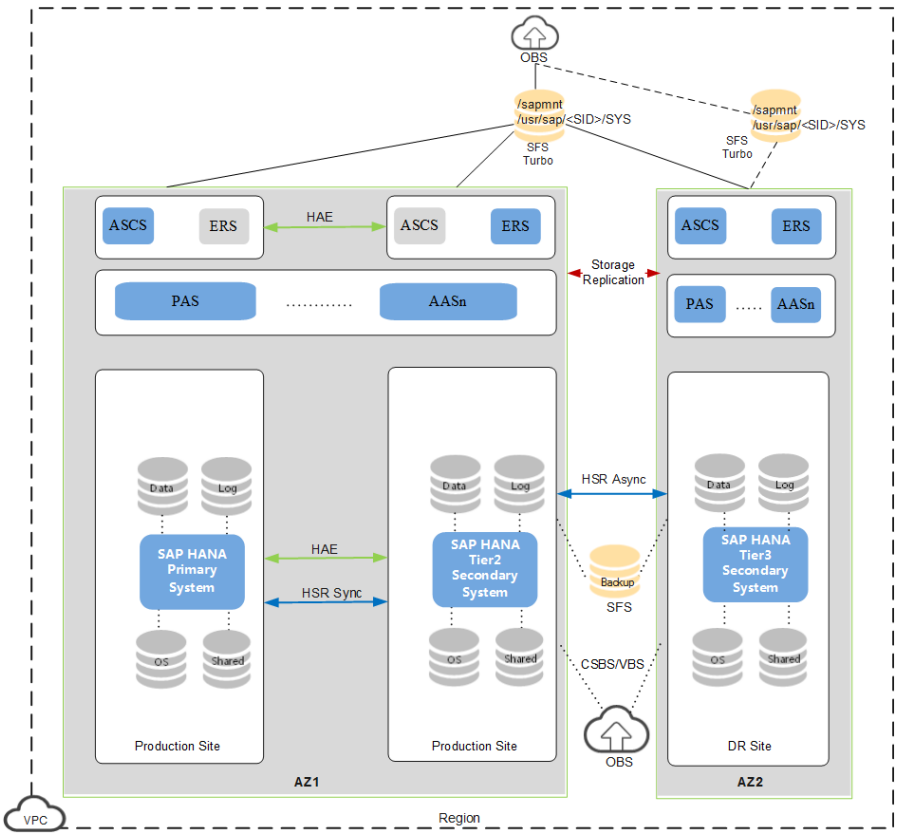

Solution Description

SAP applications are deployed in distributed mode with HA, and SAP HANA is deployed in single-node mode with HA required. AZ 1 is the production site, and AZ 2 is the DR site.

In this solution, SDRS is configured for some file systems of SAP applications, and other file systems are provided by SFS Turbo. SAP HANA adopts the multi-tier system replication (MSR) scheme. Synchronous system replication and HAE are configured in AZ 1 to meet the HA requirements of the production system. Then, asynchronous system replication is configured for the SAP HANA node in AZ 2. The OS volume, data volume, log volume, and shared volume of SAP HANA are provided by EVS in AZ 1, and the backup volume is provided by SFS, and files are shared across AZs. SAP HANA data is periodically backed up to the backup volume, or to an OBS bucket using CSBS or VBS.

If AZ 1 is faulty, SAP applications are switched over to the DR site through SDRS. The servers at the DR site become available immediately. For file systems provided by SFS Turbo, if they are successfully mounted to the DR site, the SAP applications are successfully switched to the DR site. If the original SFS Turbo file systems are faulty and cannot be recovered, you need to create SFS Turbo file systems with the same size and quantity as the originals on AZ 2, mount the new file systems to the SAP applications at the DR site, and copy the backup data on OBS to the new file systems to recover the SAP service.

SAP HANA can be restored to a specified recovery point by using the backup files in the backup volume, CSBS, or VBS.

NOTICE

After a failover, the following events may occur on a node that uses SDRS:

- **The SAP license becomes invalid.**

After SAP applications are switched to the DR site, the hardware key of the SAP applications may change because the VM IDs change. As a result, the original SAP license becomes invalid. To solve this problem, see SAP Note **2570214**.

To solve this problem, perform a DR drill when deploying the SAP system, record the hardware keys of the production and DR nodes, and apply for the required license for the two nodes on the SAP official website.

- **NIC names change.**

  Ethernet NIC names change, for example, from eth0 and eth1 to eth2 and eth3. You need to clear the udev rules and configuration files of the NIC.

  1. Run the following command to edit the udev rules files of the NICs and delete required rules based on the MAC addresses of the NICs:

     vim /etc/udev/rules.d/70-persistent-net.rules

     Figure 6-3 shows the udev rules before the deletion, and **Figure 6-4** shows the rules after the deletion.

     Figure 6-3 Before

     ```
     # net device ()
     SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="00:16:3e:2c:ee:29", ATTR{type}=="1", KERNEL=="eth*", NAME="et
     h5"

     # net device ()
     SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="00:16:3e:58:fe:7e", ATTR{type}=="1", KERNEL=="eth*", NAME="et
     h6"

     # net device ()
     SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="00:16:3e:1f:38:b0", ATTR{type}=="1", KERNEL=="eth*", NAME="et
     h7"
     ```

     Figure 6-4 After

     ```
     # net device ()
     SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="00:16:3e:2c:ee:29", ATTR{type}=="1", KERNEL=="eth*", NAME="et
     h5"

     # net device ()
     SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="00:16:3e:58:fe:7e", ATTR{type}=="1", KERNEL=="eth*", NAME="et
     h6"
     ```

  2. If the NIC configuration files have been created, run the following command to delete the NIC files:

     rm -f /etc/sysconfig/network/ifcfg-eth*x*

     ifcfg-eth*x* indicates the name of a target NIC file.

7 Related Documents

To ensure that SAP systems installed on the public cloud platform are in compliance with the SAP support requirements, you are advised to refer to the standard SAP documents and notes in **Table 7-1** and **Table 7-2**.

Table 7-1 SAP documents

Name	Link
SAP HANA Administrator Guide	**SAP HANA Administration Guide**
Setting Up SAP HANA System Replication	**Setting Up SAP HANA System Replication**
Configuring SAP HANA System Replication	**Configuring SAP HANA System Replication**
Monitoring SAP HANA System Replication	**Monitoring SAP HANA System Replication**
SAP HANA Multitier System Replication	**SAP HANA Multitier System Replication**
How to Perform System Replication for SAP HANA	**How to Perform System Replication for SAP HANA**
SAP HANA System Replication Details	**SAP HANA System Replication Details**

Table 7-2 SAP notes

Description	Note	Link
FAQ	1999880	**FAQ: SAP HANA System Replication**
Configuration description	2369981	**Required configuration steps for authentication with HANA System Replication**

Description	Note	Link
Network configuration description	1876398	**Network configuration for System Replication in SAP HANA**

A Change History

What's New	Released On
This issue is the first official release.	2019-07-12

